Measuring and Understanding Throughput of Network Topologies

Sangeetha Abdu Jyothi, Ankit Singla, P. Brighten Godfrey, Alexandra Kolla
University of Illinois at Urbana–Champaign

arXiv:1402.2531v3 [cs.NI] 9 Feb 2015

ABSTRACT

High throughput is a fundamental goal of network design and is of particular interest in data center and high performance computing networks. Although myriad network topologies have been recently proposed, a broad head-to-head comparison of these proposals is absent, and the right way to compare worst-case throughput performance turns out to be a subtle problem.

In this paper, we present a framework to benchmark the throughput of network topologies. First, we show that cut-based metrics such as bisection bandwidth are the wrong measure: they yield incorrect conclusions about throughput performance. Next, we show how to generate a near-worst-case traffic matrix for a given topology, which empirically approaches a theoretical lower bound. Finally, we employ this metric, along with other traffic matrices, to analyze the throughput performance of a variety of networks proposed for data centers and high performance computing. Our evaluation code is freely available to facilitate future reproducible work on rigorously designing and evaluating networks.

1. INTRODUCTION

Throughput is a fundamental property of communication networks: how much data can be carried across the network between desired end-points per unit time? Particularly in the application areas of data centers and high performance computing (HPC), an increase in throughput demand among compute elements has reinvigorated research on the subject, and a large number of network topologies have been proposed in the past few years to achieve high capacity at low cost [1–11].

However, there is little order to this large and ever-growing set of network topologies. We lack a broad comparison of topologies, and there is no open, public framework available for testing and comparing topology designs. The absence of a well specified benchmark complicates research on network design, making it difficult to evaluate a new design against the numerous past proposals, and difficult for industry to know which threads of research are most promising to adopt.

In fact, not only are we lacking throughput comparisons across a spectrum of topologies, we argue that the situation is worse: the community has, in many cases, been using the wrong metrics for measuring throughput. Cut-based metrics such as bisection bandwidth and sparsest cut are commonly used to estimate throughput, because minimum cuts are assumed to measure worst-case throughput [12]. However, while this is true for the case where the network carries only one flow, it does not hold for the common case of general traffic matrices. In the latter case, the cut is only an upper bound on throughput. The resulting gap between the cut-metric and the throughput leaves open the possibility that in a comparison of two topologies, the cut-metric could be larger in one topology, while the throughput could be greater in the other. Later in this paper, we shall demonstrate that this discrepancy indeed occurs in practice.

Our goal is to build a framework for accurate and consistent measurement of the throughput of network topologies, and use this framework to benchmark proposed topologies. We make the following contributions towards this goal:

(1) Bisection bandwidth is the wrong metric. We show that the commonly used cut-based metrics, bisection bandwidth and sparsest cut, are flawed predictors of performance: topology A can have a higher cut-metric than topology B while topology B has asymptotically higher throughput, even for worst-case traffic. In addition, these cut metrics are NP-complete to compute. We implement a suite of heuristics to find sparse cuts, and show that the resulting sparse cuts do in fact differ from throughput for a number of proposed topologies.

(2) Towards a throughput metric. Since cuts are the wrong metric, one can instead measure throughput directly. But throughput depends on the traffic workload presented. Therefore it is useful to test networks with worst-case traffic matrices (TMs). We develop an efficient heuristic for generating a near-worst-case TM for any given topology, and show that these near-worst-case TMs approach a theoretical lower bound of throughput in practice. One can therefore benchmark topologies using this TM, along with other workloads of interest. Finally, to make fair comparisons across proposals with differing equipment (e.g. different numbers of switches or distribution of links), we compare each topology against a random graph constructed with the same equipment, and use the throughput relative to the random graph as a consistent measure of the topology's throughput.

(3) Benchmarking proposed topologies. We evaluate a set of 10 topologies proposed for data centers and high performance computing. We find that:

- Throughput differs depending on TM, with some topologies dropping substantially in performance for near-worst-case TMs.
- The bulk of proposals, for all TMs, perform worse than networks based on random graphs [9], with the latter's advantage increasing with scale.

- Proposals for networks based on expander graphs, including Jellyfish's random topology [9], Long Hop [13], and Slim Fly [14], have nearly identical performance for uniform traffic (e.g. all-to-all). This contrasts with conclusions of [13,14]; note that [14] analyzed bisection bandwidth rather than throughput.

To the best of our knowledge, this work is the most expansive comparison of network topology throughput to date, and the only one to use an accurate and consistent method of comparison. Our evaluation framework and the set of topologies we tested are freely available [15]. We hope these tools will facilitate future work on rigorously designing and evaluating networks, and allow quick reproduction of research results.

2. METRICS FOR THROUGHPUT

In this section, we investigate the metrics currently used for estimating throughput, identify their shortcomings and propose a new metric which is more efficient and computationally tractable.

2.1 Throughput defined

Our focus in this paper is on network throughput. Furthermore, our focus is on abstract *network topology*, not on systems-level design like routing and congestion control.[1] Therefore, the metric of interest is end-to-end throughput supported by a network in a fluid-flow model with optimal routing. We next define this more precisely.

A **network** is a graph $G = (V, E)$ with capacities $c(u, v)$ for every edge $(u, v) \in E_G$. Among the nodes V are **servers**, which send and receive traffic flows, connected through non-terminal nodes called **switches**. Each server is connected to one switch, and each switch is connected to zero or more servers, and other switches. Unless otherwise specified, for switch-to-switch edges (u, v), we set $c(u, v) = 1$, while server-to-switch links have infinite capacity. This allows us to stress-test the network topology itself, rather than the servers.

A **traffic matrix (TM)** T defines the traffic demand: for any two servers v and w, $T(v, w)$ is an amount of requested flow from v to w. We assume without loss of generality that the traffic matrix is normalized so that it conforms to the "hose model": each server sends at most 1 unit of traffic and receives at most 1 unit of traffic ($\forall v$, $\sum_w T(v, w) \leq 1$ and $\sum_w T(w, v) \leq 1$).

The **throughput** of a network G with TM T is defined as the maximum value t for which $T \cdot t$ is feasible in G. That is, we seek the maximum t for which there exists a feasible multicommodity flow that routes flow $T(v, w) \cdot t$ through the network from each v to each w, subject to the link capacity and the usual flow conservation constraints. This can be formulated in a standard way as a linear program (which we omit for brevity) and is thus computable in polynomial time. If the nonzero traffic demands $T(v, w)$ have equal weight, as they will throughout this paper, this is equivalent to maximizing the minimum throughput of any of the requested end-to-end flows. This maximization is a standard problem called *maximum concurrent flow* [16].

Note that we could alternately maximize the *total* throughput of all flows. We avoid this because it would allow the network to pick and choose among the TM's flows, giving high bandwidth only to the "easy" ones (e.g., short flows that do not cross a bottleneck). The formulation above is equivalent to maximizing total throughput subject to a fairness constraint across flows.

We now have a precise definition of throughput, but it depends on the choice of traffic matrix. How can we evaluate a *topology* independent of assumptions about the traffic?

[1]There are many performance measures for networks that are not considered here, and higher-layer systems obviously affect throughput experienced by applications. But our goal is a comprehensive understanding of the fundamental goal of high-throughput network topology, which as we will see is sufficient for a single paper. At this stage, systems-level issues would be a distraction.

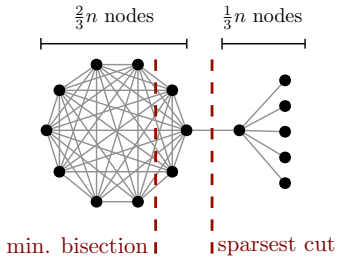

Figure 1: Bisection bandwidth fails to find the true bottleneck.

2.2 Bisection bandwidth defined

Bisection bandwidth is by far the most commonly-used attempt to provide an evaluation of a topology's performance independent of a specific TM. Since any cut in the graph upper-bounds the flow across the cut, if we find the minimum cut then we can bound the worst-case performance of the network. The intuition is that this corresponds to the unfortunate case that all communicating pairs are located on opposite sides of this cut. Now, of course, the smallest cuts might just slice off a single node, while we are interested in larger-scale bottlenecks; so the bisection bandwidth requires cuts which split the nodes into two equal-sized groups. Specifically, bisection bandwidth of a graph G is typically (see [12], p. 974) defined as

$$BB(G) = \min_{S \subseteq V, |S| = \frac{n}{2}} c(S, \bar{S}),$$

where $\bar{S}$ is the complement of S and $c(S, \bar{S})$ is the total capacity of edges crossing the bisection $(S, \bar{S})$.

2.3 Bisection bandwidth is the wrong metric

Bisection bandwidth does give some insight into the capacity of a network. It provides an upper-bound on worst-case network performance, is simple to state, and can sometimes be "eyeballed" for simple networks.

However, it has several limitations. First, it is NP-complete to compute [17]; the best approximation algorithm has a polylogarithmic approximation factor [18]. Even if one could use specialized algorithms to calculate it for structured topologies, no such methods are known for unstructured networks (such as several recently proposed for data centers [9,11,19]). This makes comparisons across networks difficult.

Second, the insistence on splitting the network in half means that bisection bandwidth may not uncover the true bottleneck. In the graph G of Figure 1, any bisection must split off at least $\frac{1}{6}n$ of the nodes in the large clique, and each

of these have $\Theta(n)$ neighbors, meaning that $BB(G) = \Theta(n^2)$ while the graph actually has a bottleneck consisting of just a single link!

This is, of course, a rather trivial problem. It may not arise in practice for highly symmetric networks typically encountered in the HPC community where bisection bandwidth has historically been used. However, it is worth mentioning here since (1) the above definition of bisection bandwidth has been used extensively, (2) the above issue means bisection bandwidth is not a sufficiently robust metric to handle general irregular and heterogeneous networks including those we consider in this paper. We also note that recently-proposed data center networks are both heterogeneous [2] and irregularly structured [9,11,19,20], so bisection bandwidth may be unsuitable for them.

Next, we will see how this particular problem with bisection bandwidth can be fixed, but how cut-based metrics (including bisection bandwidth) are subject to a more fundamental problem.

2.4 Sparsest cut defined

There are several ways to fix bisection bandwidth's overly rigid requirement of exact bisection, such as balanced partitioning [21] and sparsest cut. Our overall conclusions are not sensitive to the distinction; we use sparsest cut here.

The **uniform sparsest cut** weights the cut by the number of separated vertex pairs. That is, the uniform sparsity of a cut $S \subseteq V$ is

$$\phi(S) = \frac{c(S, \bar{S})}{|S| \cdot |\bar{S}|},$$

where $c(S, \bar{S})$ is the total capacity of edges crossing the cut. The uniform sparsest cut is the cut of sparsity $min_{S \subseteq V} \phi(S)$.

To motivate this definition, it helps to generalize it. The (nonuniform) sparsity of a cut weights the cut by the amount of *traffic demand* across the cut. That is, for a given TM T, the sparsity of a cut is

$$\phi(S, T) = \frac{c(S, \bar{S})}{T(S, \bar{S})},$$

where $T(S, \bar{S})$ is the traffic demand crossing the cut, i.e., $T(S, \bar{S}) = \sum_{v \in S, w \in \bar{S}} T(v, w)$. The **nonuniform sparsest cut** for TM T is then the cut of sparsity $min_{S \subseteq V} \phi(S, T)$. Intuitively, this cut should be the one across which it's hardest to push the traffic demand T.

Observe that the uniform sparsest cut is a special case of the nonuniform sparsest cut when we take T to be the complete traffic matrix ($T(v, w) = 1 \ \forall v, w$). Also observe that bisection bandwidth is equivalent to the uniform sparsest cut if it happens that the uniform sparsest cut has $\frac{n}{2}$ nodes on either side. The reader can assume sparsest cut refers to the uniform case unless otherwise specified.

2.5 Sparsest cut is the wrong metric

The improved cut metric will now succeed in finding the true bottleneck in Figure 1, unlike bisection bandwidth. But have we found the right metric for worst-case throughput? We argue that the answer is no, for three reasons.

(1) Sparsest cut and bisection bandwidth are not actually TM-independent, contrary to our original goal of evaluating a topology independent of traffic. As discussed above, bisection bandwidth and the uniform sparsest cut correspond to the worst cuts for the complete (all-to-all) traffic matrix, so they have a hidden implicit assumption of this particular TM.

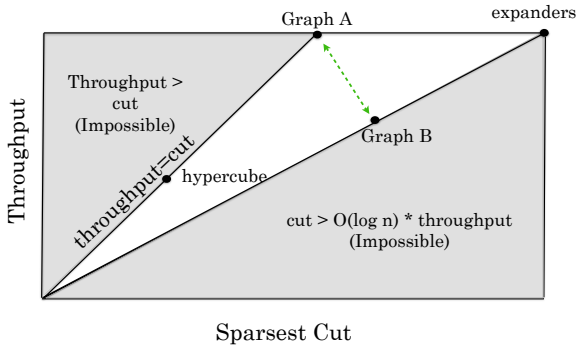

Figure 2: Sparsest cut vs. Throughput

(2) Even for a specific traffic matrix, computing sparsest cut is NP-hard for most TMs, and it is believed that there is no efficient constant factor approximation algorithm [22]. In contrast, throughput is computable in polynomial time. Sparsest cut's difficulty is both practically inconvenient and strong evidence that cut-based metrics are not computing the same physical quantity as throughput.

(3) While cuts upper-bound throughput, it is only a *loose* upper bound. This may be counterintuitive, if our intuition is guided by the well-known max-flow min-cut theorem. In a network with a single $v \rightarrow w$ traffic flow, the theorem states that the maximum $v \rightarrow w$ flow is equal to the minimum capacity over all cuts separating v and w [23, 24]. Hu [25] extended these results and showed that maximum concurrent flow is equal to the minimum cut for two-commodity flows. But this result is not valid for multi-commodity flows when the number of flows is greater than two, which is of course the case of primary interest in computer networks. Specifically, in a multi-commodity flow with uniform demands, the maximum flow throughput can be an $O(\log n)$ factor lower than the sparsest cut [26]. Hence, the maximum flow in the network is not directly captured by minimum cuts in the network.

Figure 2 depicts the relationship between sparsest cut and throughput. The flow (throughput) in the network cannot exceed the upper bound imposed by the worst-case cut. On the other hand, the cut cannot be more than a factor $O(\log n)$ greater than the flow [26]. Thus, any graph and an associated traffic matrix can be represented by a unique point in the region bounded by these limits.

While this distinction is of course well-established in the theory literature [26], we further argue here that *it can lead to incorrect decisions when evaluating networks.* Specifically, we will exhibit a pair of graphs A and B such that, as shown in Figure 2, A has higher throughput but B has higher sparsest cut. If sparsest cut is the metric used to choose a network, graph B will be wrongly assessed to have better performance than graph A, while in fact it has a factor $\Omega(\sqrt{\log n})$ worse performance!

Graph A is a clustered random graph adapted from previous work [27] with n nodes and degree $2d$. A is composed of two equal-sized clusters with $n/2$ nodes each. Each node

in a cluster is connected by degree α to nodes inside its cluster, and degree β to nodes in the other cluster, such that $\alpha + \beta = 2d$. A is sampled uniformly at random from the space of all graphs satisfying these constraints. We can pick α and β such that $\beta = \Theta(\frac{\alpha}{\log n})$. Then, as per [27] (Lemma 3), the throughput of A, T_A and its sparsest cut, Φ_A are both $\Theta(\frac{1}{n \log n})$.

Let graph G be any $2d$-regular expander on $N = \frac{n}{dp}$ nodes, where d is a constant and p is a parameter we shall adjust later. **Graph B** is constructed by replacing each edge of G with a path of length p. It is easy to see that B has n nodes. We prove in Appendix A, the following theorem.

Theorem 1. *$T_B = O(\frac{1}{np \log n})$ and $\Phi_B = \Omega(\frac{1}{np})$.*

In the above, setting $p = 1$ corresponds to the 'expanders' point in Figure 2: both A and B have the same throughput (within constant factors), but the B's cut is larger by $O(\log n)$. Increasing p creates an asymptotic separation in both the cut and the throughput such that $\Phi_A < \Phi_B$, while $T_A > T_B$. For instance, if $p = \sqrt{\log n}$, $T_B = O(\frac{1}{n(\log n)^{3/2}})$ and $\Phi_B = \Omega(\frac{1}{n\sqrt{\log n}})$. Further, if $p = \Theta(\log n)$, we can tune the constants such that $T_A > T_B \Theta(\log n)$ even though $\Phi_A < \Phi_B$.

Intuition. The reason that throughput may be smaller than sparsest cut is that in addition to being limited by bottlenecks, the network is limited by the total volume of "work" it has to accomplish within its total link capacity. That is, if the TM has equal-weight flows,

$$Throughput\,per\,flow \leq \frac{Total\,link\,capacity}{\#\,of\,flows \cdot Average\,path\,length}$$

where the total capacity is $\sum_{(i,j)\in E} c(i,j)$ and the average path length is computed over the flows in the TM. This "volumetric" upper bound may be tighter than a cut-based bound. For example, [27] observed that the cut between two halves of a random graph could be reduced significantly without affecting throughput, indicating that the throughput of a random graph is limited by the volumetric upper bound, rather than by the sparsest cut. Our Graph A corresponds to the extreme point of that process, where the cut is as small as possible without impeding throughput.

2.6 Aside: Are cut metrics useful at all?

While throughput is the focus of this paper, it is not the only graph property of interest, and cut-based metrics do capture important graph properties.

One obvious example is reliability: cut-based metrics directly capture the difficulty of partitioning a network.

A less obvious example is physical network layout. Consider two networks A, B with the same throughput, but with B having greater bisection bandwidth. Suppose we desire a physical realization of these networks grouping them into two equal-sized clusters, such that shorter, cheaper cables will be used to create links within clusters and longer, more expensive cables must be used across clusters. Since A has lower bisection bandwidth, there is a way to partition it which uses fewer long cables. In other words, in this case, B's higher bisection bandwidth is actually bad — it yields no performance benefit and corresponds to greater difficulty in physical cabling! — so we actually wish to *minimize* bisection bandwidth.

2.7 Towards a throughput metric

Having exhaused cut-based metrics, we return to the original metric of throughput defined in §2.1, and suggest a simple solution: network topologies should simply be evaluated directly in terms of throughput (via LP optimization software), for specific traffic matrices.

The key, of course, is how to choose the TM. If we can find a worst-case TM, this would fulfill the goal of evaluating topologies independent of assumptions about traffic. However, computing a worst-case TM is an unsolved, computationally non-trivial problem [28] [2]. Here, we offer practical, efficient heuristics to find a *bad-case TM* which can be used to benchmark topologies.

We begin with the complete or **all-to-all TM** T_{A2A} which for all v, w has $T_{A2A}(v, w) = \frac{1}{n}$. We observe that T_{A2A} is in fact within $2\times$ of the worst case TM. This fact is simple and known to some researchers, but at the same time, we have not seen it previously written down, so we give the statement here and proof in Appendix B.

Theorem 2. *Let G be any graph. If T_{A2A} is feasible in G with throughput t, then any hose model traffic matrix is feasible in G with throughput $\geq t/2$.*

Can we get closer to the worst case TM? In our experience, TMs with a smaller number of "elephant" flows are more difficult to route than TMs with a large number of small flows, like T_{A2A}. This suggests a **random matching TM** in which we have only one outgoing flow and one incoming flow per server, chosen uniform-randomly among servers. We can further stress the network topology by decreasing the number of servers attached to each switch, so there are even fewer flows originating from each switch. These TMs actually lack the factor of 2 near-worst-case guarantee of T_{A2A}, but in practice we have found they consistently result in lower throughput than T_{A2A}.

Can we get *even closer* to the worst-case TM? Intuitively, the all-to-all and random matching TMs will tend to find sparse cuts, but only have average-length paths. Drawing on the intuition that throughput decreases with average flow path length, we seek to produce traffic matrices that force the use of long paths. To do this, given a network G we compute all-pairs shortest paths and create a complete bipartite graph H, whose nodes represent all sources and destinations in G, and for which the weight of edge $v \to w$ is the length of the shortest $v \to w$ path in G. We then find (via another LP optimization) the maximum weight matching in H. The resulting matching corresponds to the pairing of servers which maximizes average flow path length, assuming flow is routed on shortest paths between each pair. We call this a **longest matching TM**.

Figure 3 shows the resulting throughput of these TMs in three topologies: hypercubes, random regular graphs, and fat trees [1]. Note that a comparison of absolute throughput across topologies is meaningless since these are not built with the same "equipment" (node degree, number of servers, etc.); the meaningful comparison here is of the various TMs within each topology. In all cases, A2A traffic has the high-

[2]Our problem corresponds to the separation problem of the minimum-cost robust network design in section 3 of [28]. This problem is shown to be hard for single-source hose model. However, the complexity is unknown for hose model with multiple sources which is the scenario under consideration.

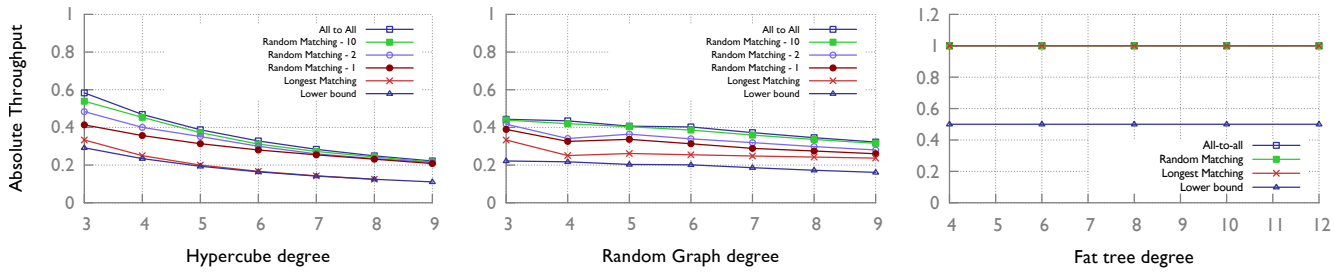

Figure 3: Throughput resulting from several different traffic matrices in three topologies

est throughput; throughput decreases or does not change as we move to a random matching TM with 10 servers per switch, and progressively decreases as the number of servers per switch is decreased to 1 under random matching, and finally to the longest matching TM. We also plot the lower bound given by Theorem 2: $T_{A2A}/2$.

These three topologies were chosen to illustrate cases when our approach is most helpful, somewhat helpful, and least helpful at finding near-worst-case TMs. In the **hypercube**, the longest matching TM is extremely close to the worst-case performance. To see why, note that the longest paths have length d in a d-dimensional hypercube, twice as long as the mean path length; and the hypercube has $n \cdot d$ unidirectional links. The total flow in the network will thus be $(\#\,flows \cdot average\,flow\,path\,length) = n \cdot d$. Thus, all links will be perfectly utilized if the flow can be routed, which empirically it is. In the **random graph**, there is less variation in end-to-end path length, but across our experiments the longest matching is always within 1.5× of the provable lower bound (and may be even closer to the true lower bound, since Theorem 2 may not be tight). In the **fat tree**, which is here built as a three-level nonblocking topology, there is essentially no variation in path length since asymptotically nearly all paths go three hops up to the core switches and three hops down to the destination. Here, our TMs are no worse than all-to-all, and the simple $T_{A2A}/2$ lower bound is off by a factor of 2 from the true worst case (which is throughput of 1 as this is a nonblocking topology).

2.8 Summary and implications

Cut-based metrics can be safely abandoned for the purposes of topology performance evaluation. They are not in fact independent of assumptions on the traffic matrix; they are NP-complete to compute; and they are not always an accurate measure of worst-case throughput, which as we showed can lead to incorrect performance evaluations. These facts call into question work which has optimized networks based on bisection bandwidth [29] and work which compared networks on "equal footing" by attempting to approximately equalize their bisection bandwidth [30].

Having surmounted that moment of catharsis, we actually obtain some relief: directly evaluating throughput with particular TMs using LP optimization is both more accurate and more tractable than cut-based metrics. In choosing a TM to evaluate, both "average case" and near-worst-case TMs are reasonable choices. All-to-all, random permutation, and measurements from real-world applications may offer reasonable average-case benchmarks. For near-worst-case traffic, we developed a practical heuristic that often succeeds in substantially worsening the TM compared with all-to-all.

Note that measuring throughput directly is not in itself a novel idea: numerous papers have evaluated particular topologies on particular TMs. In contrast, our study provides a rigorous analysis of throughput vs. cut-metrics, a way to generate near-worst-case traffic for any given topology, and — in what follows — an experimental evalution of both our metrics and of a large number of proposed topologies in those metrics.

3. EXPERIMENTAL EVALUATION

Having established a framework for evaluating throughput, in this section we evaluate the framework experimentally: Are cut-metrics indeed worse predictors of performance? When measuring throughput directly, how close do we come to worst-case traffic? Then, we employ our framework for a comparison of commonly-used and proposed data center and HPC network topologies. But first, we present our experimental methodology.

3.1 Methodology

3.1.1 *Computing cut-metrics*

Before we can evaluate whether cuts predict throughput, we need to deal with the fact that the cut-metrics themselves are NP-complete to compute. Bisection bandwidth considers all cuts that divide the network in half, and sparsest cut considers all possible cuts; both are exponentially large sets.We focus our evaluation on sparsest-cut, the more accurate indicator of the flow in the network.

To deal with the computational complexity, we implemented a set of heuristics that find potentially bad cuts, apply all of them, and finally use the sparsest cut found by any heuristic. These heuristics are: (1) Brute force, which is feasible only for networks of up to around 20 nodes, after which we have the code simply output the sparsest of around 100,000 considered cuts. (2) Cutting every single node and (3) every pair of nodes, to catch cases when sparse cuts are near the edge of the network. (4) Cutting expanding regions of the network, defined by nodes within distance k from each node, for each relevant value of k. (5) An eigenvector-based technique which is known to be within a constant factor of the actual sparsest cut. Details of these heuristics are in Appendix C. We thus made every effort to use the best available means for calculating cut metrics, lest we dismiss them casually. We believe this is the most thorough empirical evaluation of cut metrics that has appeared for data center networks.

3.1.2 Computing throughput

Throughput is computed as a solution to a linear program whose objective is to maximize the minimum flow across all flow demands, as explained in §2. We use the Gurobi [31] solver to compute throughput. Throughput depends on the traffic workload provided to the network.

3.1.3 Traffic workload

We evaluate three traffic matrix families given in §2.7 – all to all, random matching and our hard-to-route longest matching. In addition, we need to specify where the traffic endpoints are located. In structured networks with restrictions on server-locations, servers are added at the locations prescribed by the models. For example, in fat-trees, servers are attached only to the lowest layer. For all other networks, we either add 1 server or k servers to each switch.

3.1.4 Topologies

Our evaluation uses 10 families of computer networks, plus a collection of "natural" networks.

BCube [3], designed for modular data center networks, differs from other structured computer networks in its server-centric design. $BCube_k$ consists of $(k+1)$ levels with each level containing n^k n-port switches. Each server in $BCube_k$ has degree $(k+1)$ and is connected to one switch in each of the $(k+1)$ levels. $(k+1) * n^k$ switches in BCube supports exactly n^{k+1} hosts.

DCell [5] is also designed for data center networks, and like BCube is server-centric. It is constructed as a recursive design starting with a building block where a switch is connected to a small number of servers. $DCell_0$ has n servers connected to a switch. $DCell_1$ is constructed from $n+1$ copies of $DCell_0$. The number of servers supported by $DCell_k$ is $t_k = t_{k-1} \cdot (t_{k-1} + 1)$.

Dragonfly [32] combines a group of routers to form a virtual router with a higher effective radix, thereby reducing the network diameter. This topology is primarily intended for HPC environments. A Dragonfly with a routers per group, where each router has h connections to other groups and p connections to servers, can support $ap(ah+1)$ servers.

Fat Tree [33] based designs are used in modern data centers and HPC. A three-level fat-tree constructed from degree-k switches consists of $\frac{5}{4}k^2$ switches and supports $\frac{k^3}{4}$ hosts.

Flattened butterfly [34] is used in on-chip networks. A k-ary n-fly topology consists of k^{n-1} switches.

Hypercubes [35] are used in HPC and computer architecture to interconnect processors. A d-dimensional hypercube consists of 2^d nodes and has a bisection bandwidth of 2^{d-1}.

HyperX [36] introduced a technique for designing HPC topologies that searches through combinations of flattened butterfly and hypercube topologies to find the optimal configuration to support a given number of servers with desired bisection bandwidth.

Jellyfish [9,27] is a uniform-random degree-specified graph. It was proposed for data center networks to enable flexible design and expansion, and to achieve higher throughput than structured designs. Since Jellyfish can be constructed for any size and node degree distribution, we use it as a reference point for direct comparison with other proposals.

Long Hop networks [13] are designed for data centers with the objective of maximizing the bisection bandwidth of the topology. This is achieved by translating best-known linear error correction codes to optimal designs in a family of Cayley networks. Networks of size 2^d can be constructed with node degrees varying from d to $2^d - 1$.

Slim Fly [14] is designed for HPC environments with the objective of reducing network diameter. The average path length of Slim Fly is equal to the theoretical lower bound. However, these networks can be constructed only for a limited number of node degrees.

Natural networks. For evaluating the cut-based metrics in a wider variety of environments, we also consider 66 non-computer networks – food webs, social networks, and more. See [15] for details on these topologies.

3.2 Evaluation of the metrics

In this section, we present the experimental results comparing cut-based metrics and the new throughput metric. We show that:

- For the majority of networks, bisection bandwidth and sparsest-cut do not predict throughput accurately, in the sense that worst-case estimated cut differs from the computed throughput by a large factor.
- Throughput under different traffic matrices on the same network consistently obeys $T_{A2A} \geq T_{RM(k)} \geq T_{RM(1)} \geq T_{LM} \geq T_{LB}$, i.e., all-to-all is consistently the easiest TM in this set, followed by random matchings, longest matchings, and the theoretically established lower bound.
- The heuristic for near-worst-case traffic, the longest matching TM, is a significantly more difficult TM than A2A and RM and often approaches the lower bound.

3.2.1 Do sparse cuts predict throughput?

We generated multiple networks from each of our topology families (with varying parameters such as size and node degree), computed throughput with the longest matching TM, and found sparse cuts using the heuristics of §3.1.1 with the same longest matching TM. Note that we will use the term **sparse cut** to refer to the sparsest cut that was found by any of the heuristics (as opposed to the true sparsest cut).

Figure 4 shows the results. Sparse cuts differ from throughput substantially, with up to a 3× discrepancy. In only a small number of cases, the cut equals throughput. Table 1 shows this in more detail: comparing column 2 and the last column, we see that cuts accurately predicted throughput in less than 15% of the tested networks only. The majority of these were from our set of natural networks; the difference between cut and flow is more pronounced on computer networks. For example, Figure 4(b) shows that although HyperX networks of different sizes have approximately same flow value, they differ widely when sparse cuts are considered. This shows that estimation of worst-case throughput performance of networks using cut-based metrics can lead to erroneous results.

How well did our sparse cut heuristics perform? Table 1 shows how often each estimator found the sparse cut. More than one technique may find the sparse cut, hence the sum may not equal the total number of networks. Brute-force

Topology family	# networks with throughput = estimated cut	Sparsest cut estimator					
		Brute force	1-node	2-node	Expanding regions	Eigenvector	Total
BCube	2	2	0	0	3	7	7
DCell	2	2	0	0	2	3	4
Dragonfly	0	2	0	0	0	2	4
Fat tree	8	8	8	8	8	8	8
Flattened butterfly	5	6	0	1	0	5	8
Hypercube	3	3	0	0	1	6	7
HyperX	1	1	0	0	1	10	11
Jellyfish	3	0	0	0	2	349	350
LongHop	9	45	0	0	1	66	110
SlimFly	1	1	0	0	0	5	6
Natural networks	48	18	21	11	34	38	66
Total	82	88	29	20	52	499	581

Table 1: Estimated sparsest cuts: Do they match throughput, and which estimators produced those cuts?

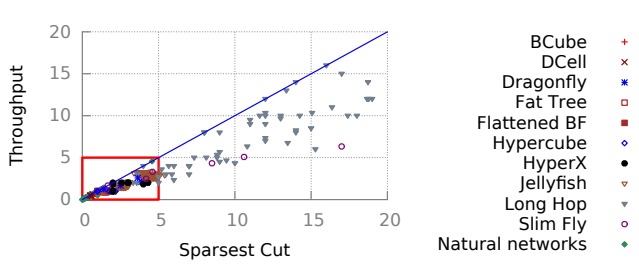

(a) Throughput vs. cut for all graphs

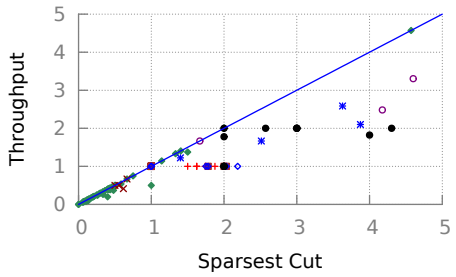

(b) Throughput vs. cut for selected graphs

Figure 4: Throughput vs. cut. (Comparison is valid for individual networks, not *across* networks, since they have different numbers of nodes and degree distributions.)

computation was helpful in finding 15% of the sparse cuts. Cuts involving one or two nodes and contiguous regions of the graph also found the sparse cut in a small fraction of networks (less than 10% each). The majority of such networks are the natural networks, which are often denser in the core and sparser in the edges. Sparse connectivity at the edges lead to bottlenecks at the edge which are revealed by cuts involving one or two nodes. Fat tree is another interesting case where every heuristic's cuts yield the accurate flow value. Overall, the eigenvector-based approximation found the largest number of sparse cuts (86%), but it is known not to be a tight approximation [37], and the full collection of heuristics did improve on it in a nontrivial fraction of cases.

Even with this suite of techniques, our calculation of sparse cuts is not guaranteed to find the *sparsest* cut, which could be closer to throughput. But even if that is the case, we can conclude that sparse cuts *that are feasible to compute in practice* are poor predictors of throughput (in addition to being computationally demanding). This complements our theoretical results (§2.5).

3.2.2 *Can we find near-worst-case traffic?*

In this section, we evaluate our proposed throughput metric — in particular, how closely the traffic matrices of §2.7 approach the worst case. We compare representative samples from each family of network under four types of traffic matrix: all to all (A2A), random matching with 5 servers per switch, random matching with 1 server per switch, and longest matching.

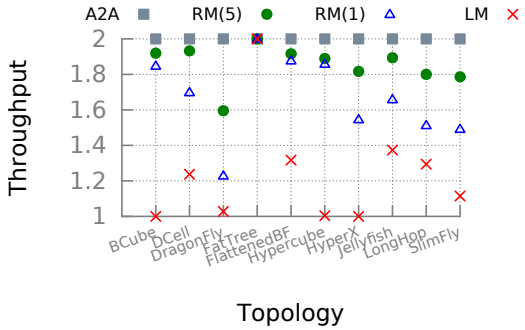

Figure 5: Comparison of throughput under different traffic matrices normalized with respect to theoretical lower bound

Figure 5 shows the throughput values normalized so that the theoretical lower bound on throughput is 1, and therefore A2A's throughput is 2. For all networks, $T_{A2A} \geq T_{RM(5)} \geq T_{RM(1)} \geq T_{LM} \geq 1$, matching the intuition discussed in §2.7. (As in Figure 4, throughput comparisons are valid across TMs for a particular network, not across networks. The networks used in this plot support approximately 250 servers each. However, the exact number of switches, links and servers varies across networks.)

Our longest matching TM is successful in matching the

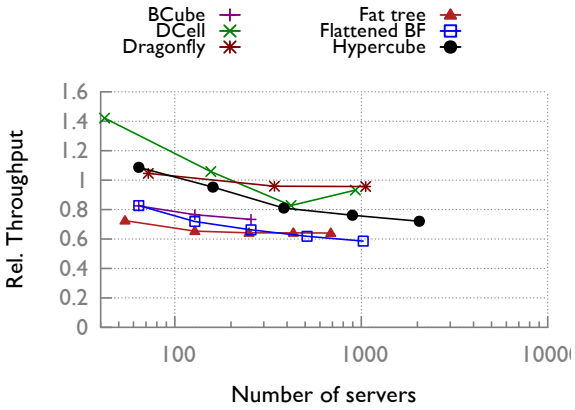

(a) All to All TM

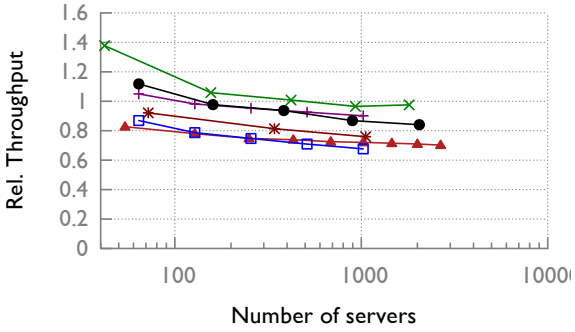

(b) Random Matching TM

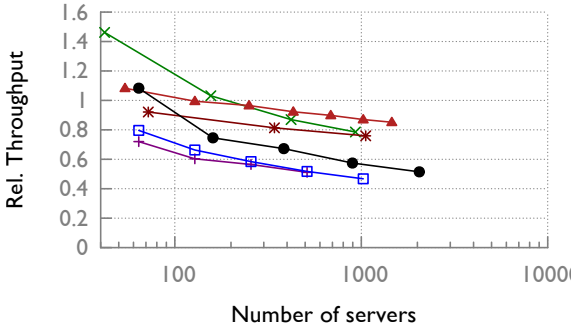

(c) Longest Matching TM

Figure 6: Comparison of TMs on topologies

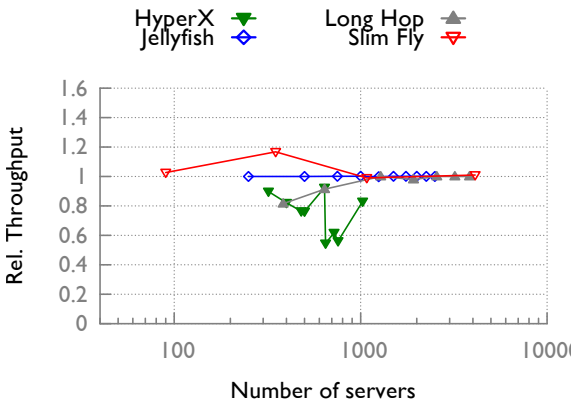

(a) All to All TM

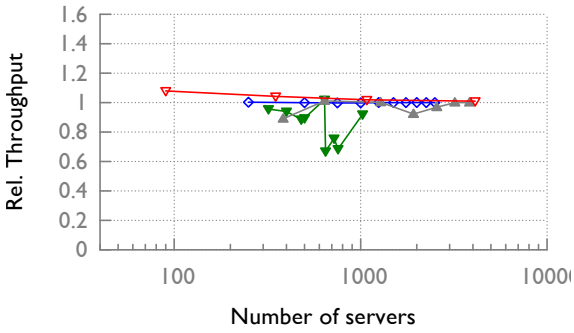

(b) Random Matching TM

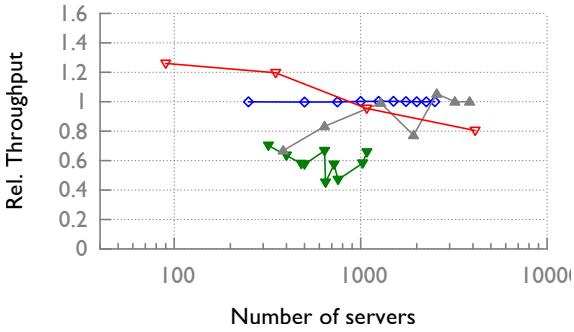

(c) Longest Matching TM

Figure 7: Comparison of TMs on more topologies

lower bound for BCube, Hypercube, HyperX, and (nearly) Dragonfly. In all other families except fat trees, the traffic under longest matching is significantly closer to the lower bound than with the other TMs. In fat trees, throughput under A2A and longest matching are equal. However, this is a characteristic of the network and not a shortcoming of the metric. In fat trees, it can be easily verified that the normalized traffic is the same under all traffic matrices. In short, these results show that throughput measurement using longest matching is a more accurate estimate of worst-case throughput performance than cut-based approximations, in addition to being substantially easier to compute.

3.3 Evaluation of topologies

Although the longest matching traffic matrix can give near worst-case performance results with limited computational effort, there is one more piece of the puzzle to allow comparison of networks. Networks may be built with different equipment – with a wide variation in number of switches and links. The raw throughput value does not account for this difference in network resources, and most proposed topologies can only be built with particular discrete numbers of servers, switches, and links, which inconveniently do not match.

Fortunately, uniform-random graphs as in [9] can be constructed for any size and specified degree distribution. Hence, random graphs serve as a convenient benchmark for easy comparison of network topologies. Our high-level approach to evaluating a network is to: (i) compute the network's throughput; (ii) build a random graph with precisely the

same equipment, i.e., the same number of nodes each with the same number of links as the corresponding node in the original graph, (iii) compute the throughput of this same-equipment random graph under the same TM; (iv) normalize the network's throughput with the random graph's throughput for comparison against other networks. This normalized value is referred to as **relative throughput**. Unless otherwise stated, each data-point is an average across 10 iterations, and all error bars are 95% two-sided confidence intervals. Minimal variations lead to narrow error bars in networks of size greater than 100.

We evaluate the performance of 10 types of computer networks[3] under all three traffic matrices: all to all, random matching with one server and longest matching. Figures 6 and 7 show the results, divided among two figures for visual clarity. Most topologies are in Figure 6, while Figure 7 shows expander-like graphs (Jellyfish, Slim Fly and Long Hop) and HyperX which has widely varying trends.

Overall, performance varies substantially, by around 1.6× with A2A traffic and more than 2× with longest matching. Furthermore, for the networks of Figure 6, which topology "wins" depends on the traffic matrix. For example, Dragonfly has high relative throughput under all-to-all traffic but under the longest matching TM, fat trees achieve the highest throughput at the largest scale. However, in all TMs, the Jellyfish random graph achieves consistently highest performance (relative throughput = 1 by definition), with Long Hop also matching that performance at the largest scale. In comparison with random graphs, performance degrades for most networks with increasing size. We discuss each network's performance in more detail below. Unless otherwise stated, reported metrics correspond to the largest network tested in each family.

BCube [3] has a server-centric design with a single server connected to multiple switches. In order to emulate this behavior we replace each server with a switch connected to a single server. This is equivalent to the BCube topology. Random graphs are constructed corresponding to this BCube-equivalent topology. The 2-ary BCube achieves 73% of the random graph's throughput under all to all traffic and 90% under random matching, but only 51% under longest matching. Performance of k-ary BCube under random matching is shown in Figure 8 with throughput degrading for all values of k with increase in network size.

DCell [5] Similar to BCube, servers in DCell are replaced by equivalent switches during comparison. 5-ary DCell achieves 93% of the random graph's throughput under all to all, 97% under random matching, and 79% under longest matching. Performance of DCell constructed from switches of different sizes are given in Figure 8. We observe that the relative throughput performance degrades for all switch sizes as the network size increases.

Dragonfly [32] networks constructed with low-radix switches of degrees varying from 4 to 8 have throughput performance under all-to-all traffic equal that of random graphs, i.e., a relative throughput approximately equal to 1. However, the performance degrades to 76% under random matching and further down to 72% under longest matching for the network constructed from switches of degree 8.

Fat Tree [33] constructed from switches of degree 14 can achieve only 65% of throughput of random graph under all to all matching. Interestingly, a fat tree achieves 73% of a random graph's throughput under random matching, and 89% with longest matching. While the fat-tree's non-blocking design is overkill for the random matching, compared to networks other than the random graph, it is advantageous for the longest matching. Note that the fat-tree's *absolute* performance does not improve when moving to the more difficult longest matching TM; instead, its *relative* throughput is higher because longest matching degrades performance more in the random graph than in the fat-tree.

Flattened butterfly [34] At the largest scale, the 2-ary flattened butterfly which can support 1024 servers achieves only 59% of the random graph's throughput under all-to-all, but the performance improves to 71% under random matching. However, the worst-case performance under longest matching drops to 47% which is the worst performance across all topologies tested. Beyond a large enough scale, the larger the k for a k-ary butterfly, the better its performance as shown in Figure 8.

Hypercubes [35] achieve 72% of the random graph's throughput under all-to-all traffic, 84% under random matching, and 51% under longest matching.

HyperX [36] Unlike other topologies, the performance of HyperX is irregular. The variations arise from the choice of underlying topology: Given switch radix, number of servers and desired bisection bandwidth, HyperX attempts to find the least cost topology constructed from the building blocks – hypercube and flattened butterfly. Even a slight variation in one of the parameters can lead to a significant difference in HyperX construction and throughput.

Results are plotted for networks with bisection bandwidth of 0.4, i.e., in the worst-case the network can support flow from 40% of the hosts without congestion. In the worst-case, the HyperX network that can support 648 hosts achieves a relative throughput of 55% under all-to-all traffic. For the same network, relative performance improves to 67% under random matching but further degrades to 45% under longest matching. We investigate the performance of HyperX networks with different bisection bandwidths under longest matching TM in Figure 9 and observe that the performance varies widely with network size under all values of bisection. Moreover, high bisection does not guarantee high performance. This further shows that using a cut-based metric such as bisection bandwidth as a cornerstone for network design can lead to poor or unpredictable performance.

Jellyfish [9] is our benchmark, and its (average) relative performance is 100% under all workloads by definition. Performance depends on randomness in the construction, but at larger sizes, performance differences between Jellyfish instances are minimal ($\approx 1\%$).

Long Hop networks [13] In Figure 10, we show that the relative throughput of Long Hop networks approaches 1 at large sizes under longest matching TM. Similar trends are observed under both all-to-all and random matching workloads. The paper [13] claimed high performance (in terms of bisection bandwidth) with substantially less equipment

[3]We realize there are always more topologies available to test, so we have made our benchmarking tools available [15] for other researchers to easily add to these comparisons.

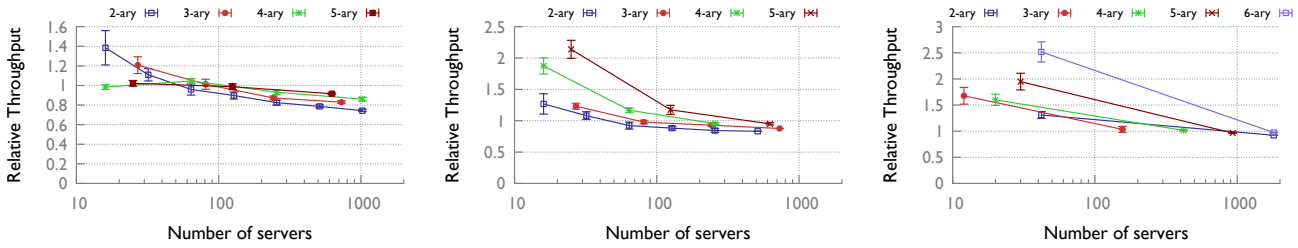

Figure 8: Relative throughput with random matching TM and varying topological parameters, in (a) Flattened butterfly, (b) BCube, and (c) DCell

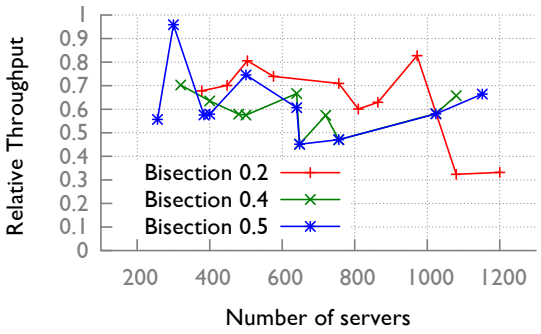

Figure 9: HyperX Relative Throughput under Longest Matching with different Bisection

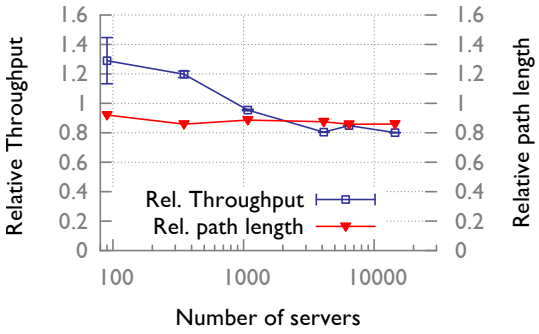

Figure 11: Slim Fly Relative Throughput under Longest Matching TM

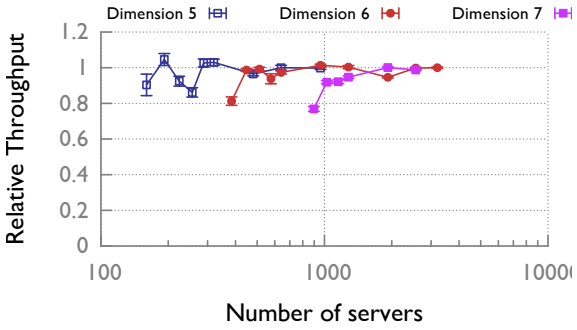

Figure 10: Long Hop network relative throughput under longest matching TM

than past designs, but we find that while Long Hop networks do have high performance, they are no better than random graphs.

Slim Fly [14] It was noted [14] that Slim Fly's key advantage is path length. We observe in Figure 11 that the topology indeed has extremely short paths – about 85-90% relative to the random graph, which is itself near-optimal – but this does not translate to higher throughput. Slim Fly's performance is essentially identical to random graphs under all-to-all and random matching TMs with a relative performance of 101% using both. Relative throughput under the longest matching TM decreases with scale, dropping to 80% at the largest size tested. Hence, throughput performance cannot be predicted solely based on the average path length.

4. RELATED WORK

The literature on network topology design is large and growing quickly, with a number of designs having been proposed in the past few years [1–11]. However, each of these research proposals only makes a comparison with one or two other past proposals, with no standard benchmarks for the comparison. There has been no independent, rigorous evaluation of a large number of topologies.

The only significant work in the space is from Popa et. al. [38]. They assess 4 topologies to determine the one that incurs least expense while achieving a target level of performance under a specific workload (all-to-all traffic). Their attempts to equalize performance required careful calibration, and approximations still had to be made. Accounting for the different costs of building different topologies is also an inexact process. We sidestep that issue by using the random graph as a normalizer: instead of attempting to match performance, for each topology, we build a random graph with *identical* equipment, and then compare throughput performance of the topology with that of the random graph. Each topology is compared to the random graph and thus the problem of massaging structured designs into roughly equivalent configurations is alleviated. This also makes it easy for others to use our tools, and to test arbitrary workloads. Apart from comparing topologies, our work also argues the superiority of flow-metrics to cuts.

Other work on comparing topologies is more focused on reliability and cuts in the topology [39].

Several researchers have used bisection bandwidth and sparsest cut as proxies for throughput performance. Fur-

ther, the usage of these two terms is not consistent across the literature. For instance, REWIRE [11] explicitly optimizes its topology designs for high sparsest cut, although it refers to the standard sparsest cut metric as bisection bandwidth. Tomic [29] builds topologies with the objective of maximizing bisection bandwidth (in a particular class of graphs). Webb et. al [40] use bisection bandwidth to pick virtual topologies over the physical topology. An interesting point of note is that they consider all-to-all traffic "a worst-case communication scenario", while our results (Figure 5) show that other traffic patterns can be significantly worse. PAST [30] tests 3 data center network proposals with the same sparsest cut (while referring to it as bisection bandwidth). Although it is not stated, the authors presumably used approximate methods because even approximating sparsest cut is believed to be NP-Hard [22]. While that does not imply that it could not be efficiently computed for these particular graphs, to the best of our knowledge, no sparsest-cut computation procedures are known for any of the networks tested in PAST. Further, PAST finds that the throughput performance of topologies with the same sparsest cut is different in packet-level simulations, raising questions about the usefulness of such a comparison; one must either build topologies of the same cost and compare them on throughput (as we do), or build topologies with the same performance and compare cost (as Popa et. al [38] do). These observations underscore the community's lack of clarity on the relationship between bisection bandwidth, sparsest cut, and throughput. A significant component of our work in this paper tackles this subject.

5. CONCLUSION

Although high throughput is a core goal of network design, measuring throughput is a subtle problem. We have shown that one common way of measuring throughput performance, via bisection and sparsest cut, produces inaccurate predictions of performance. We have put forward an improved benchmark for network capacity, including a near-worst-case traffic matrix, and have used this metric to compare a large number of proposed network topologies, revealing new perspective on performance of several proposals. Particularly when designing networks for environments where traffic patterns may be unknown or variable and performance requirements are stringent, we believe evaluation under the near-worst-case longest matching TM will be useful.

Our findings also raise interesting directions for future work.

- Our longest-matching traffic heuristically produces near-worst case performance, but does not always match the lower bound. Is there an efficient method to produce even-worse-case traffic for any given topology, or provably approximate it within less than 2×?
- We observed that low average path length does not guarantee high throughput. How is throughput correlated with other fundamental graph properties? High correlations could suggest improvement in network design, while small correlation would further position throughput as an independent topological property of interest.

Our topology benchmarking tools and the topology datasets are all freely available [15].

6. REFERENCES

[1] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A scalable, commodity data center network architecture. In *SIGCOMM*, 2008.

[2] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A. Maltz, Parveen Patel, and Sudipta Sengupta. Vl2: A scalable and flexible data center network. In *SIGCOMM*, 2009.

[3] C. Guo, G. Lu, D. Li, H. Wu, X. Zhang, Y. Shi, C. Tian, Y. Zhang, and S. Lu. BCube: a high performance, server-centric network architecture for modular data centers. *SIGCOMM Comput. Commun. Rev.*, 39(4):63–74, August 2009.

[4] Radhika Niranjan Mysore, Andreas Pamboris, Nathan Farrington, Nelson Huang, Pardis Miri, Sivasankar Radhakrishnan, Vikram Subramanya, and Amin Vahdat. Portland: A scalable fault-tolerant layer 2 data center network fabric. In *SIGCOMM*, 2009.

[5] Chuanxiong Guo, Haitao Wu, Kun Tan, Lei Shi, Yongguang Zhang, and Songwu Lu. Dcell: A scalable and fault-tolerant network structure for data centers. In *SIGCOMM*, 2008.

[6] Guohui Wang, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, T. S. Eugene Ng, Michael Kozuch, and Michael Ryan. c-through: Part-time optics in data centers. In *SIGCOMM*, 2010.

[7] Nathan Farrington, George Porter, Sivasankar Radhakrishnan, Hamid Hajabdolali Bazzaz, Vikram Subramanya, Yeshaiahu Fainman, George Papen, and Amin Vahdat. Helios: A hybrid electrical/optical switch architecture for modular data centers. In *SIGCOMM*, 2010.

[8] Ankit Singla, Atul Singh, Kishore Ramachandran, Lei Xu, and Yueping Zhang. Proteus: a topology malleable data center network. In *HotNets*, 2010.

[9] Ankit Singla, Chi-Yao Hong, Lucian Popa, and P. Brighten Godfrey. Jellyfish: Networking data centers randomly. In *NSDI*, 2012.

[10] Andrew R. Curtis, S. Keshav, and Alejandro Lopez-Ortiz. LEGUP: using heterogeneity to reduce the cost of data center network upgrades. In *CoNEXT*, 2010.

[11] Andrew R. Curtis, Tommy Carpenter, Mustafa Elsheikh, Alejandro Lopez-Ortiz, and S. Keshav. Rewire: An optimization-based framework for data center network design. Technical report, August 2011.

[12] D. Padua. *Encyclopedia of Parallel Computing*. Number v. 4 in Springer reference. Springer, 2011.

[13] Ratko V. Tomic. Optimal networks from error correcting codes. In *ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, 2013.

[14] M. Besta and T. Hoefler. Slim Fly: A Cost Effective Low-Diameter Network Topology. In *IEEE/ACM International Conference on High Performance Computing, Networking, Storage and Analysis (SC14)*, Nov. 2014.

[15] Obscured for review. Our topology benchmark tool, 2013.

[16] F. Shahrokhi and D.W. Matula. The maximum concurrent flow problem. *Journal of the ACM*,

37(2):318–334, 1990.

[17] Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. 1979.

[18] Uriel Feige and Robert Krauthgamer. A polylogarithmic approximation of the minimum bisection. *SIAM Journal on Computing*, 31(4):1090–1118, 2002.

[19] Sonja Filiposka and Carlos Juiz. Community-based complex cloud data center. *Physica A: Statistical Mechanics and its Applications*, 419(0):356 – 372, 2015.

[20] Ji-Yong Shin, Bernard Wong, and Emin GÃijn Sirer. Small-world datacenters. *ACM SOCC*, 2011.

[21] Konstantin Andreev and Harald Racke. Balanced graph partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.

[22] Shuchi Chawla, Robert Krauthgamer, Ravi Kumar, Yuval Rabani, and D Sivakumar. On the hardness of approximating multicut and sparsest-cut. *computational complexity*, 2006.

[23] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399 –404, 1956.

[24] P. Elias, A. Feinstein, and C. Shannon. A note on the maximum flow through a network. *Information Theory, IEEE Transactions on*, 2(4):117–119, Dec 1956.

[25] T. C. Hu. Multicommodity network flows. *Oper. Res.*, 11(3):344 – 360, 1963.

[26] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, November 1999.

[27] Ankit Singla, P. Brighten Godfrey, and Alexandra Kolla. High throughput data center topology design. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2014.

[28] Chandra Chekuri. Routing and network design with robustness to changing or uncertain traffic demands. *SIGACT News*, 38(3):106–129, September 2007.

[29] Ratko V. Tomic. Optimal networks from error correcting codes. In *ANCS*, 2013.

[30] Brent Stephens, Alan Cox, Wes Felter, Colin Dixon, and John Carter. Past: Scalable ethernet for data centers. In *CoNEXT*, 2012.

[31] Gurobi Optimization Inc. Gurobi optimizer reference manual. `http://www.gurobi.com`, 2013.

[32] John Kim, William J. Dally, Steve Scott, and Dennis Abts. Technology-Driven, Highly-Scalable Dragonfly Topology. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, ISCA '08, pages 77–88, 2008.

[33] Charles E. Leiserson. Fat-trees: universal networks for hardware-efficient supercomputing. *IEEE Trans. Comput.*, 34(10):892–901, October 1985.

[34] John Kim, William J. Dally, and Dennis Abts. Flattened butterfly: a cost-efficient topology for high-radix networks. *SIGARCH Comput. Archit. News*, 35(2):126–137, June 2007.

[35] `http://hdl.handle.net/2027.42/27522`.

[36] Jung Ho Ahn, Nathan Binkert, Al Davis, Moray McLaren, and Robert S. Schreiber. Hyperx: Topology, routing, and packaging of efficient large-scale networks. In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis*, SC '09, 2009.

[37] F.R.K. Chung. Laplacians of graphs and cheeger's inequalities. In *Combinatorics, Paul ErdÃűs is Eighty, Vol. 2*, pages 157–172. Janos Bolyai Mathematical Society, Budapest, 1996.

[38] Lucian Popa, Sylvia Ratnasamy, Gianluca Iannaccone, Arvind Krishnamurthy, and Ion Stoica. A cost comparison of datacenter network architectures. In *CoNEXT*, 2010.

[39] David S. Lee and Jeffrey L. Kalb. Network topology analysis. Technical report, 2008.

[40] Kevin C Webb, Alex C Snoeren, and Kenneth Yocum. Topology switching for data center networks. In *HotICE*, 2011.

[41] Moshe Babaioff and John Chuang. On the optimality and interconnection of valiant load-balancing networks. In *INFOCOM*, 2007.

APPENDIX

A. PROOF OF THEOREM 1

We revisit the *maximum concurrent flow* problem, based on which we defined throughput in §2: Given a network $G = (V, E_G)$ with capacities $c(u, v)$ for every edge $(u, v) \in E_G$, and a collection (not necessarily disjoint) of pairs $(s_i, t_i), i = 1, \ldots, k$ each having a unit flow demand, we are interested in maximizing the minimum flow. Instead of the traffic matrix (TM) formulation of §2, for the following discussion, it will be convenient to think of the pairs of vertices that require flow between them as defining a demand graph, $H = (V, E_H)$. Thus, given G and H, we want the maximum throughput. As we noted in §2, this problem can be formulated as a standard linear program, and is thus computable in polynomial time.

We are interested in comparing our suggested throughput metric with sparsest cut. We first prove the following theorem.

Theorem 3. *The dual of the linear program for computing throughput is a linear programming relaxation for sparsest cut.*

Proof. We shall use a formulation of the throughput linear program that involves an exponential number of variables but for which is easier to derive the dual. We denote by $P_{s,t}$ the set of all paths from s to t in G and we introduce a variable x_p for each path $p \in P_{s,t}$, for each $(s, t) \in E_H$, corresponding to how many units of flow from s to t are routed through path p.

$$\begin{aligned}
\max \quad & y \\
\text{subject to} \quad & \sum_{p \in P_{s,t}} x_p \geq y && \forall (s,t) \in E_H, \\
& \sum_{p:(u,v)\in p} x_p \leq c(u,v) && \forall (u,v) \in E_G \\
& x_p \geq 0 && \forall p \\
& y \geq 0.
\end{aligned}$$

The dual of the above linear program will have one variable $w(s, t)$ for each $(s, t) \in E_H$ and one variable $z(u, v)$ for each $(u, v) \in E_G$.

$$\begin{aligned}
\min \quad & \sum_{u,v} z(u,v)c(u,v) \\
\text{subject to} \quad & \sum_{(s,t)\in E_H} w(s,t) \geq 1 \\
& \sum_{(u,v)\in p} z(u,v) \geq w(s,t) && \forall (s,t) \in E_H, p \in P_{s,t} \\
& w(s,t) \geq 0 && \forall (s,t) \in E_H \\
& z(u,v) \geq 0 && \forall (u,v) \in E_G.
\end{aligned}$$

It is not hard to realize that in an optimal solution, without loss of generality, $w(s, t)$ is the length of the shortest path from s to t in the graph weighted by the $z(u, v)$. We can also observe that in an optimal solution we have $\sum w(s, t) = 1$. These remarks imply that the above dual is equivalent to the following program, where we introduce a variable $l(x, y)$ for every pair or vertices in $E_G \cup E_H$.

$$\begin{aligned}
\min \quad & \sum_{u,v} l(u,v)c(u,v) \\
\text{subject to} \quad & \sum_{(s,t)\in E_H} l(s,t) = 1 \\
& \sum_{(u,v)\in p} l(u,v) \geq l(s,t) && \forall (s,t) \in E_H, p \in P_{s,t} \\
& l(u,v) \geq 0 && \forall (u,v) \in E_G \cup E_H
\end{aligned}$$

The constraints $\sum_{(u,v)\in p} l(u, v) \geq l(s, t)$ can be equivalently restated as triangle inequalities. This means that we require $l(u, v)$ to be a metric over V. These observations give us one more alternative formulation:

$$\min_{l(\cdot,\cdot) \text{ metric}} \frac{\sum_{(u,v)\in E_G} c(u,v) \cdot l(u,v)}{\sum_{(s,t)\in E_H} l(s,t)} \tag{1}$$

We can finally see that the above formulation is a linear programming relaxation for a cut problem. More specifically, the sparsest cut problem is asking to find a cut S that minimizes the ratio

$$\frac{\sum_{(u,v)\in E_G \text{ cut by S}} c(u,v)}{|\text{ edges} \in E_H \text{ cut by S}|} \tag{2}$$

This is equivalent to minimizing ratio (1) over ℓ_1 metrics only.

If we take E_H to be the complete graph (corresponding to all-to-all demands), we get the standard sparsest cut definition:

$$\frac{\sum_{(u,v)\in E_G \text{ cut by S}} c(u,v)}{|S||\bar{S}|} \tag{3}$$

□

Before we prove Theorem 1 from §2, we shall demonstrate the following claim.

Claim 4. *If G is a d-regular expander graph on N nodes and H is the complete graph, the value of the linear program for throughput is $(\frac{d \log d}{N \log N})$. The value of the sparsest cut is $\Omega(\frac{d}{N})$.*

Proof. Let us denote by T the optimal value of expression (1). Note that this is the optimal value of the dual for the linear program for throughput, therefore equal to the optimal throughput. By taking $l(\cdot, cdot)$ to be the shortest path metric on G, we calculate:

$$T \leq \frac{\sum_{(i,j)\in E_G} l(i,j)}{\sum_{i,j\in V} l(i,j)} \leq \frac{d/2.|V|}{\Theta(N^2 \frac{\log N}{\log d})} \leq O(\frac{d \log d}{N^2 \log N}) \tag{4}$$

Here, the first inequality follows from the fact that for d-regular graphs, each node can reach no more than d^i nodes in i hops. This means that given a vertex v, there exist $\Theta(N)$ nodes with distance at least $\frac{\log N}{\log d}$ from it, which means that the total distance between all pairs of nodes is $\Theta(N^2 \frac{\log N}{\log d})$.

Let Φ denote the minimum value of ratio (3) for G. Since G is an expander, this ratio is

$$\Phi \geq \Omega(\frac{d \cdot |S|}{|S||V-S|}) = \Omega(\frac{d}{N}) \quad (5)$$

□

Theorem 1. *Let graph G be any $2d$-regular expander on $N = \frac{n}{dp}$ nodes, where d is a constant and p is a free parameter. Let graph B be constructed by replacing each edge of G with a path of length p. Then, B has throughput $T_B = O(\frac{1}{np\log n})$ and sparsest cut $\Phi_B = \Omega(\frac{1}{np})$.*

Proof. Let (S_1, S_2) be the sparsest cut in B. Let $(S_1{}', S_2{}')$ be the corresponding cut in G. Namely, if an edge was cut in B by (S_1, S_2) that belonged to a path p_e then $(S_1{}', S_2{}')$ cuts e. Let Φ_B be the value of the cut (S_1, S_2) in A and Φ_G the value of $(S_1{}', S_2{}')$ in G. Then

$$\Phi_B = \frac{E(S_1, S_2)}{|S_1||S_2|} = \frac{E(S_1{}', S_2{}')}{|S_1||S_2|} \geq \frac{E(S_1{}', S_2{}')}{p \cdot |S_1'|p \cdot |S_2'|} \geq \frac{\Phi_G}{p^2}$$

by equation (5) we have $\Phi_G \geq \Omega(\frac{1}{N}) = \Omega(\frac{p}{n})$ which gives us

$$\Phi_B \geq \Omega(\frac{1}{np})$$

On the other hand, let T_B be the value of the throughput of B. We follow a similar reasoning as we did in equation (4).

$$\begin{aligned} T_B \leq \frac{\sum_{(i,j)\in E_G} l(i,j)}{\sum_{i,j\in V} l(i,j)} &\leq \frac{Ndp}{\Theta((Np)^2 p \log N)} \\ &\leq O(\frac{1}{Np^2 \log N}) = O(\frac{1}{np \log n}) \end{aligned} \quad (6)$$

□

B. PROOF OF THEOREM 2

Proof. T_{A2A} has demand $\frac{1}{n}$ on each flow, so the largest feasible multicommodity flow routing of T_{A2A} in G has capacity $\frac{t}{n}$ allocated to each flow. Let C be a graph representing this routing, i.e., a complete digraph with capacity $\frac{t}{n}$ on each link. Systems-oriented readers may find it useful to think of C as an overlay network implemented with reserved bandwidth in G. In other words, to prove the theorem, it is sufficient to show that taking T to be any hose-model traffic matrix, $T \cdot t/2$ is feasible in C.

For this, we use a two-hop routing scheme analogous to Valiant load balancing [41]. Consider any traffic demand $v \rightsquigarrow w$. In the first step, we split this demand flow into n equal parts, routing flow $\frac{1}{n} \cdot T(v,w) \cdot t/2$ from v to every node in the network, along the direct links (or the zero-hop path when the target is v itself). In the second step, the traffic arriving at each node is sent along at most one link to its final destination.

We now have to show that this routing is feasible in C. Consider any link $i \to j$. This link will carry a fraction $\frac{1}{n}$ of all the traffic originated by i, and a fraction $\frac{1}{n}$ of all the traffic destined to j. Because T is a hose model TM, each node originates and sinks a total of ≤ 1 unit of traffic; and since we are actually attempting to route the scaled traffic matrix $T \cdot t/2$, each node originates and sinks a total of $\leq t/2$ units of traffic. Therefore, link i carries a total of

$$\frac{t}{2} \cdot \frac{1}{n} + \frac{t}{2} \cdot \frac{1}{n} = \frac{t}{n},$$

which is the available capacity on each link of C and is hence feasible. □

C. MEASURING CUTS

We employ several heuristics for estimating sparsest cut.

Brute-force computation Brute force computation of sparsest cut is computationally intensive since it considers all possible cuts in the network (2^{n-1} cuts in a network with n nodes). In addition to bandwidth, the number of flows traversing each cut has to be estimated which adds further overhead in the computation of sparsest cut.

Due to the computational complexity, brute force evaluation of sparsest cut is possible only for networks of size less than 20. However, we perform limited brute-force computation on all networks by capping the computation at $10,000$ cuts.

One-node cuts Designed computer networks as well as naturally occurring networks tend to be denser at the core and sparse at the edges. When the core has high capacity, it is likely that the worst-case cut occurs at the edges. Hence, this heuristic considers all cuts with only a single node in a subset formed by the cut. There exists n cuts with a single node. This is a very small fraction of the total 2^{n-1} cuts.

Two-node cuts $\frac{n*(n-1)}{2}$ cuts with two nodes in a subset also reveal the limited connectivity at the edges of the network.

Expanding cuts It is likely that the network is clustered,i.e., it contains two or more highly connected components connected by a few links. Subsets of all possible combinations of contiguous nodes in the network can be very large. We optimize our search to a subspace of this category of cuts. Starting from each node, we consider cuts which include nodes within a distance k from the node. When $k = 0$, the cut involves only the originating node and is equivalent to the single node case discussed before. When $k = 1$, all nodes within distance 1 from the node are considered – the node and its neighbors. k is incremented until the entire graph is covered. If d is the diameter of the network, the number of cuts considered is less than or equal to $n * d$.

Eigenvector based optimizations Eigenvector corresponding to the second smallest eigenvalue of the normalized Laplacian of a graph can give a set of n cuts, the worst of which is within a constant factor from the actual cut [37]. The nodes of the graph are sorted in the ascending order corresponding to their value in the second eigenvector [37]. We sweep this vector of sorted nodes to obtain the n cuts.